arXiv:1410.1358v3 [math.GT] 30 Apr 2016

The pseudo-Anosov and conjugacy problems are in $\mathbf{NP} \cap \mathbf{co\text{-}NP}$

Mark C. Bell
University of Illinois
mcbell@illinois.edu

May 3, 2016

Abstract

For a fixed marked surface S, we construct polynomial bounds on the periodic and preperiodic lengths of the maximal splitting sequences of a projectively invariant measured train track.

We give two consequences of these bounds. Firstly, that the problem of deciding whether a mapping class is pseudo-Anosov lies in **NP**. This is dual to the previously known result that the pseudo-Anosov problem is in **co-NP**. Secondly, that the problem of deciding whether two mapping classes are conjugate lies in **co-NP**. Similarly, this is the dual to the previously known result that the conjugacy problem is in **NP**.

As usual, in both cases we immediately obtain exponential time solutions to these problems. A version of these algorithms have been implemented as part of `flipper`.

keywords. mapping class group; pseudo-Anosov; conjugacy; $\mathbf{NP} \cap \mathbf{co\text{-}NP}$.

Mathematics Subject Classification (2010): 37E30, 68W40, 57-04

1 Introduction

Fix S to be a marked surface with at least one marked point (or at least three when the surface has no genus).

Fix X to be a finite generating set of $\mathrm{Mod}^+(S)$, the *mapping class group* of S (relative to the set of marked points). Let X^* denote the set of all *words* that can be made using the elements of X as letters. We identify a word $h = h_1 \cdots h_k \in X^*$ with the mapping class $h_1 \circ \cdots \circ h_k$ and denote its *length* by $\ell(h) := k$.

Let $\mathcal{ML}(S)$ denote the space of *measured laminations* on S [19, Section 1.7] [7, Section 3]. Recall that a mapping class $h \in \mathrm{Mod}^+(S)$ is *pseudo-Anosov* [11, Exposè 12] [7, Page 95] if there is a measured lamination $\mathcal{L} \in \mathcal{ML}(S)$ which is:

- *projectively invariant*, that is, $h(\mathcal{L}) = \lambda \cdot \mathcal{L}$ for some $\lambda \in \mathbb{R}$, and
- *filling*, that is, it assigns positive measure to every essential, simple closed curve.

We say that a word is *pseudo-Anosov* if its corresponding mapping class is.

Problem 1.1 (The Pseudo-Anosov Problem)**.** Given a word $h \in X^*$, decide whether h is pseudo-Anosov.

If h is not pseudo-Anosov then by the Nielsen–Thurston classification [10, Chapter 13] it is either *periodic* or *reducible*. There is a polynomial time solution to the word problem for $\mathrm{Mod}^+(S)$ [10, Theorem 4.2] and an upper bound on the order of finite order mapping classes [10, Theorem 7.5]. Together these give a polynomial time algorithm to determine whether h is *periodic*. On the other hand, there are several proofs that if h is reducible then there is a reducing curve whose complexity is at most $O(\mathrm{poly}(\ell(h)))$ [4, Theorem 3.3] [13, Theorem 1.1]. Such a curve acts as a certificate of reducibility which can be verified in polynomial time. This shows that the pseudo-Anosov problem is in **co-NP** [4, Corollary 3.5].

In Section 3 we prove the dual result. Namely, if $h \in X^*$ is pseudo-Anosov then there is a description of a measured lamination $\mathcal{L}$ such that it can be checked that $\mathcal{L}$ is projectively invariant and filling in $O(\mathrm{poly}(\ell(h)))$ time. That is, if a word is pseudo-Anosov then there is a small proof of this, or equivalently:

Corollary 3.20. *The pseudo-Anosov problem lies in* ***NP****.* □

In the second half of this paper, we consider the *conjugacy problem*. Here we write $g \equiv h$ if two words $g, h \in X^*$ represent the same mapping class.

Problem 1.2 (The Conjugacy Problem)**.** Given words $g, h \in X^*$, decide whether g and h represent *conjugate* mapping classes. That is, decide whether there is there a word $f \in X^*$ such that $fgf^{-1} \equiv h$.

Theorem B of [22] states that there there is a constant $K = K(S, X)$ such that if two words $g, h \in X^*$ represent conjugate mapping classes then there is a word $f \in X^*$ such that

$$h \equiv fgf^{-1} \quad \text{and} \quad \ell(f) \leq K \cdot (\ell(g) + \ell(h)).$$

Such a word f acts as a certificate that g and h are conjugate and, as it is sufficiently small, this can be verified in polynomial time [16]. This shows that the conjugacy problem is in **NP**.

In Section 4 we use the results developed in Section 2 to again show the dual result. Namely, if a pair of words do not represent conjugate mapping classes then there is a short proof of this, or equivalently:

Corollary 4.15. *The conjugacy problem lies in* ***co-NP****.* □

Finally, in Section 5 we discuss some important implementation details of these algorithms.

1.1 Model of computation

The operations that we will describe will involve manipulating reasonably large numbers. We will take care to include the difficultly of performing these in our analysis of the algorithm.

Definition 1.3. An integer x is *k–bounded* if it has at most k digits, that is, if it can be represented by $O(k)$ bits.

As part of our model of computation, we will assume that if x and y are k–bounded then:

- $\text{sign}(x)$ can be computed in $O(1)$ operations,
- $x \pm y$ is $(k+1)$–bounded and can be computed in $O(k)$ operations, and
- xy is $2k$–bounded and can be computed in $O(k^2)$ operations.

However, for ease of argument, we will assume a model of computation in which there is no cost associated to accessing variables.

1.2 Models of $\mathcal{ML}(S)$ and $\text{Mod}^+(S)$

Fix $\mathcal{T}$ to be an (ideal) triangulation of S with (ordered) edges $e_1, \ldots, e_\zeta$. The triangulation $\mathcal{T}$ provides a coordinate system on $\mathcal{ML}(S)$ where the lamination $\mathcal{L} \in \mathcal{ML}(S)$ is represented by the *edge vector*:

$$\mathcal{T}(\mathcal{L}) := \begin{pmatrix} \mathcal{L}(e_1) \\ \vdots \\ \mathcal{L}(e_\zeta) \end{pmatrix} \in \mathbb{R}^\zeta_{\geq 0}$$

Although distinct laminations have distinct coordinates, not all coordinates correspond to measured laminations. In fact $v_1, \ldots, v_\zeta$ corresponds to a measured lamination if and only if:

- $\sum v_i > 0$,
- for each triangle a, b, c it satisfies the triangle inequality $v_a + v_b \geq v_c$, and
- for each vertex v there is an incident triangle a, b, c such that $a \cap b = v$ and $v_a + v_b = v_c$.

We express this requirement as a linear programming problem:

Lemma 1.4. *For each triangulation $\mathcal{T}$, there are $O(1)$–bounded $\zeta \times 3\zeta$ matrices $F_1, \ldots, F_k$ such that $v \in \mathbb{R}^\zeta_{\geq 0}$ is in the image of $\mathcal{T}(\cdot)$ if and only if $v \neq 0$ and*

$$F_i \cdot v \geq 0$$

for some i. □

Let G denote the *graph of (labelled) triangulations* of S. This is a graph whose vertices are triangulations of S and two are connected via an edge of length one if and only if they differ by a single *flip*, as shown in Figure 1, or a *reordering* or the edges. As G and $\text{Mod}^+(S)$ are quasi-isometric [6, Proposition 8.19], we may think of a word $h \in X^*$ as a path p in G from $\mathcal{T}$ to $h(\mathcal{T})$. Moreover we may choose this path such that $\ell(p) \in O(\ell(h))$.

The edge vector of a lamination is clearly well behaved under a reordering but it is also well behaved under a flip.

Proposition 1.5 ([15, Page 30])**.** *Suppose that $\mathcal{L}$ is a measured lamination and e is a flippable edge of a triangulation $\mathcal{T}$ as shown in Figure 1 then*

$$\mathcal{L}(e') = \max(\mathcal{L}(a) + \mathcal{L}(c), \mathcal{L}(b) + \mathcal{L}(d)) - \mathcal{L}(e).$$ □

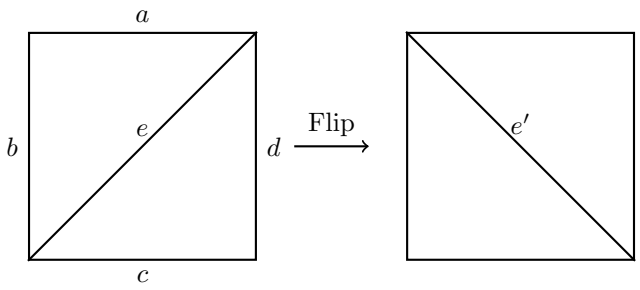

Figure 1: Flipping an edge of a triangulation.

Thus there is a piecewise-linear function which transforms $\mathcal{T}$ coordinates to $\mathcal{T}'$ coordinates. We express this piecewise-linear function using two collections of matrices $\{A_i\}$ and $\{B_i\}$.

Lemma 1.6. *Suppose that p is a path from $\mathcal{T}$ to $\mathcal{T}'$. There are matrices $\{A_i\}$ and $\{B_i\}$, as described in [4, Section 2.2], such that:*

1. *Each A_i and B_i is $\ell(p)$–bounded.*
2. *Each B_i has $O(\ell(p))$ rows.*
3. *For each measured lamination $\mathcal{L} \in \mathcal{ML}(S)$ we have that $B_i \cdot \mathcal{T}(\mathcal{L}) \geq 0$ for some i.*
4. *For each measured lamination $\mathcal{L} \in \mathcal{ML}(S)$, if $B_i \cdot \mathcal{T}(\mathcal{L}) \geq 0$ then $\mathcal{T}'(\mathcal{L}) = A_i \cdot \mathcal{T}(\mathcal{L})$.* □

Remark 1.7. Suppose that p is a path from $\mathcal{T}$ to $\mathcal{T}'$ and that $\{A_i\}$ and $\{B_i\}$ are the matrices of Lemma 1.6. If v is a vector in which each entry has at most d digits then by considering each flip in turn we can find B_i such that $B_i \cdot v \geq 0$ in at most $O(d\ell(p) + \ell(p)^2)$ operations. Moreover each entry of $A_i \cdot v$ has at most $d + \ell(p)$ digits.

2 Train tracks

We begin with a discussion of train tracks [19]. These will be crucial for determining if a measured lamination is filling. To ease notation, throughout this section we will assume that $h \in X^*$ is a fixed word and that $\mathcal{L} \in \mathcal{ML}(S)$ is a measured lamination such that:

$$h(\mathcal{L}) = \lambda \cdot \mathcal{L}.$$

Definition 2.1 ([19, Section 1.1])**.** A *measured train track (representing $\mathcal{L}$)* is a pair $T = (\tau, \mu)$ consisting of:

- a *train track* τ, that is, a trivalent graph on S (whose vertices we refer to as *switches* and edges we refer to as *branches*) with a well defined tangency at each switch such that there is no switch with all branches emanating from it in the same direction and no complementary region of τ is a nullogon, monogon, bigon, once-marked nullogon or annulus, and

- a *transverse measure* μ such that there is a smooth map $\phi\colon S \to S$, isotopic to the identity, such that:
  - $\phi(\mathcal{L}) = \tau$,
  - $\phi|_{\mathcal{L}}\colon \mathcal{L} \to \tau$ is a submersion, and
  - $\mu = \mathcal{L} \circ \phi^{-1}$

Definition 2.2. If $T = (\tau, \mu)$ is a measured train track then we may *split* along along one of its branches e to obtain a new measured train track $T' = (\tau', \mu')$ as shown in Figure 2 [19, Page 119].

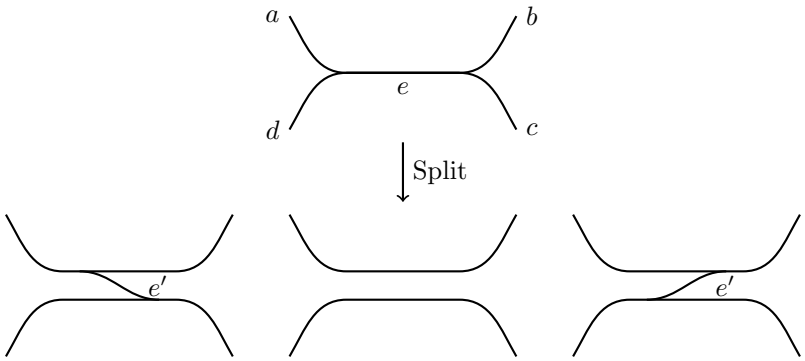

When $\mu(a) > \mu(b)$. When $\mu(a) = \mu(b)$. When $\mu(a) < \mu(b)$.

Figure 2: The possibilities for splitting a branch e. In either case $\mu(e') = |\mu(a) - \mu(b)| = |\mu(c) - \mu(d)|$ [19, Figure 2.1.2].

There are many different measured train tracks that represent $\mathcal{L}$. Again we represent these via a simplicial graph $\mathbb{G}$, the *graph of measured train tracks (representing $\mathcal{L}$)*. This is a graph whose vertices are measured train tracks (representing $\mathcal{L}$) and there is an (unoriented) edge from T to T' if and only if T' can be obtained by splitting some collection of the branches of T, each of which has the same transverse measure. This graph is connected [19, Theorem 2.8.5] and we write $d(T, T')$ for the distance between two measured train tracks when each edge of this graph is assigned length one.

2.1 Maximal splittings and the axis

Definition 2.3 ([1, Section 2])**.** The *maximal splitting* of a measured train track T is the measured train track $s(T)$ obtained by simultaneously splitting all branches of maximal measure.

Lemma 2.4. *If T and T' are measured train tracks (representing $\mathcal{L}$) then*

$$d(s(T), s(T')) \leq d(T, T').$$

That is, s is a non-expansive map.

Proof. Begin by considering the case in which $d(T, T') = 1$. Without loss of generality, there is a subset B of the branches of T such that by splitting the branches in B we obtain T' and each branch in B has the same transverse measure. Let M be the subset of the branches of T with maximal transverse measure. There are three possible cases to consider:

1. If $B = M$ then $s(T) = T'$ and so $d(s(T), s(T')) = d(T', s(T')) \leq 1$.
2. If $B \subset M$ then the branches in $M - B$ also appear in T'. It is these branches that are split when maximally splitting T' and so $s(T) = s(T')$. Hence $d(s(T), s(T')) = 0$.
3. If $B \not\subseteq M$ then the branches in B also appear in $s(T)$. Splitting these branches results in $s(T')$ as M and B lie in disjoint open sets on S. Hence $d(s(T), s(T')) = 1$.

In any case $d(s(T), s(T')) \leq 1$ and so the result holds by induction on $d(T, T')$. □

Agol showed that the sequence of train tracks obtained by repeatedly performing maximal splittings is eventually periodic (up to rescaling).

Theorem 2.5 ([1, Theorem 3.5])**.** *If T is a measured train track then there exists $m, n \in \mathbb{N}$ such that $s^{m+n}(T) = \widehat{h}(s^n(T))$ where*

$$\widehat{h}(\tau, \mu) := (h(\tau), \mu/\lambda). \qquad \square$$

We refer to the smallest such m and n as the *periodic* and *preperiodic lengths* of T respectively. We note that m depends only on $\mathcal{L}$ and is independent of T.

Definition 2.6. The *axis* of $\mathcal{L}$ is the bi-infinite sequence of measured train tracks $A = A(\mathcal{L}) := \{\mathbf{T}_i\}_{i=-\infty}^{\infty}$ such that

$$s(\mathbf{T}_i) = \mathbf{T}_{i+1} \quad \text{and} \quad \widehat{h}(\mathbf{T}_i) = \mathbf{T}_{i+m}.$$

The measured train tracks on the axis are useful as you can determine if $\mathcal{L}$ is filling purely from the combinatorics of their underlying train tracks.

Definition 2.7. A measured train track $T = (\tau, \mu)$ is *filling* if every complementary region of τ is either a disk or a once-marked disk.

Lemma 2.8. *The measured lamination $\mathcal{L}$ is filling if and only if $\mathbf{T}_i$ is.*

Proof. If $\mathbf{T}_i = (\tau, \mu)$ is not filling then there is an essential, simple closed curve γ in the complement of τ. Therefore $\mathcal{L}(\gamma) = 0$ and so $\mathcal{L}$ is not filling.

Conversely, if $\mathcal{L}$ is not filling then there is an essential, simple closed curve γ such that $\mathcal{L}(\gamma) = 0$. There is a measured train track $T = (\tau, \mu)$, representing $\mathcal{L}$, such that τ and $N(\gamma)$ are disjoint. By Theorem 2.5 there are j and k such that $\widehat{h}^k(s^j(T)) = \mathbf{T}_i$. As T is not filling and this is preserved by both maximal splittings and homeomorphisms, $\mathbf{T}_i$ is not filling either. □

2.2 Getting to the axis

We use Lemma 2.8 to decide whether $\mathcal{L}$ is filling by constructing a measured train track on the axis A. To do this we give an upper bound the periodic and preperiodic lengths of T.

Theorem 2.9. *The periodic length of T is at most $d(T, \widehat{h}(T))$.*

Proof. Suppose that

$$p := T_0, T_1, \ldots, T_k$$

is a *path* from T to $\widehat{h}(T)$, that is, a sequence such that $d(T_i, T_{i+1}) = 1$. Let $T_i^j := s^j(T_i)$. Note that for each j:

- $\widehat{h}(T_0^j) = T_k^j$, and
- $p_j := T_0^j, \ldots, T_k^j$ is also a path as s is non-expansive.

When j is sufficiently large, by Theorem 2.5 each T_i^j lies on the axis A and therefore so does p_j. The endpoints of p_j must be $\mathbf{T}_{k'}$ and $\mathbf{T}_{k'+m}$ and so we conclude that $m \leq k$.

As this holds for every such path p, we have that the periodic length of T is at most $d(T, \widehat{h}(T))$. □

To bound the preperiodic length, the following theorem is key. In particular, note that the constant depends only on the underlying surface and is independent of the chosen mapping class h and measured lamination $\mathcal{L}$.

Theorem 2.10 ([22, Theorem B])**.** *There is a constant $K = K(S)$ such that if T and T' are measured train tracks then there is a measured train track T'' in the orbit of T' under h such that*

$$d(T, T'') \leq K(d(T, \widehat{h}(T)) + d(T', \widehat{h}(T'))). \qquad \square$$

Agol also showed that for any measured train track T (representing $\mathcal{L}$), if we repeated perform maximal splittings then eventually every branch of T split [1, Lemma 2.1]. For train tracks on the axis of $\mathcal{L}$ we can give an explicit upper bound on the number of maximal splittings needed.

Lemma 2.11. *If $\mathbf{T}_i$ is a train track on the axis of $\mathcal{L}$ then every branch of $\mathbf{T}_i$ must be split within $3\zeta m$ maximal splittings.*

Proof. Let B_k be the set of branches of $\mathbf{T}_i$ that are split within km maximal splittings. Note that $B_{k+1} = B_k \cup \widehat{h}(B_k)$. Therefore if B_k does not contain of all branches of $\mathbf{T}_i$ and $|B_{k+1}| = |B_k|$ then there would be a branch that is never split, which cannot happen [1, Lemma 2.1]. Hence either B_k consists of all branches of $\mathbf{T}_i$ or $|B_{k+1}| > |B_k|$. As $\mathbf{T}_i$ has at most 3ζ branches, the latter case cannot occur when $k \geq 3\zeta$. Thus every branch of $\mathbf{T}_i$ must be split within $3\zeta m$ maximal splittings. □

We note that by this result every branch of $\mathbf{T}_i$ must become a branch of maximal transverse measure within $3\zeta m$ maximal splittings.

Theorem 2.12. *The preperiodic length of T is at most $6\zeta K d(T, \widehat{h}(T))^2$.*

Proof. We claim that if $\mathbf{T}_i$ is a measured train track on A then the preperiodic length of a measured train track T is at most $3\zeta m d(T, \mathbf{T}_i)$. We prove this by induction on the distance. There is nothing to show in the base case when $d(T, \mathbf{T}_i) = 0$ so suppose that the claim is true whenever $d(T, \mathbf{T}_i) < k$. Now if $d(T, \mathbf{T}_i) = k$ then let T' be such that

$$d(T', T) = 1 \quad \text{and} \quad d(T', \mathbf{T}_i) = k - 1$$

as shown in Figure 3. Then by induction the preperiodic length of T' is at most $3\zeta m(k-1)$ and so $\mathbf{T}_j := s^{3\zeta m(k-1)}(T')$ is on the axis A. For ease of notation let $T'' := s^{3\zeta m(k-1)}(T)$, then as s is non-expansive (Lemma 2.4)

$$d(\mathbf{T}_j, T'') \leq 1.$$

There are now three possibilities that can occur.

The first possibility is that $d(\mathbf{T}_j, T'') = 0$ and so $T'' = \mathbf{T}_j$.

The second possibility is that $d(\mathbf{T}_j, T'') = 1$ and there is a subset B of the branches of $\mathbf{T}_j$ such that every branch in B has the same transverse measure and splitting along these yields T''. Now consider $s^p(T'')$ and $s^p(\mathbf{T}_j)$ where, by Lemma 2.11, we choose $p < 3\zeta m$ such that the branches in B appear in $s^p(\mathbf{T}_j)$ with maximal measure. If there are no other branches of maximal measure then maximally splitting $s^p(\mathbf{T}_j)$ results in $s^p(T'')$. Hence $s^p(T'') = \mathbf{T}_{j+p+1}$. Otherwise maximally splitting $s^p(\mathbf{T}_j)$ factors through $s^p(T'')$ and so $s^{p+1}(T'') = \mathbf{T}_{j+p+1}$.

The third possibility is that $d(\mathbf{T}_j, T'') = 1$ and there is a subset B of the branches of T'' such that splitting along these yields $\mathbf{T}_j$. In this case let B' be new the branches that are added when the branches in B are split. Again consider $s^p(T'')$ and $s^p(\mathbf{T}_j)$ where, by Lemma 2.11, we choose $p < 3\zeta m$ such that the branches in B' appear in $s^p(\mathbf{T}_j)$ with maximal measure. Now in order for $d(s^{p+1}(\mathbf{T}_j), s^{p+1}(T'')) \leq 1$ we must have already split the branches in B. Hence $s^{p+1}(T'')$ is must be either $\mathbf{T}_{j+p+1}$ or $\mathbf{T}_{j+p}$.

In any case, it follows that $s^{3\zeta mk}(T) = s^{3\zeta m}(T'')$ is either:

$$\mathbf{T}_{j+3\zeta m-1}, \quad \mathbf{T}_{j+3\zeta m} \quad \text{or} \quad \mathbf{T}_{j+3\zeta m+1}$$

and so is on the axis A. Hence the claim is true.

Now by Theorem 2.10 there is a $\mathbf{T}_i$ on the axis A such that

$$d(T, \mathbf{T}_i) \leq K(d(T, \widehat{h}(T)) + d(\mathbf{T}_i, \widehat{h}(\mathbf{T}_i))) = K(d(T, \widehat{h}(T)) + m).$$

Therefore by the previous claim and Theorem 2.9 the preperiodic length of T is at most

$$3\zeta m K(d(T, \widehat{h}(T)) + m) \leq 6\zeta K d(T, \widehat{h}(T))^2. \qquad \square$$

Combining this with Lemma 2.8 we obtain:

Corollary 2.13. *The measured lamination $\mathcal{L}$ is filling if and only if $s^t(T)$ is filling, where $t := 6\zeta K d(T, \widehat{h}(T))^2$.* $\square$

Finally, we note that a triangulation $\mathcal{T}$ gives rise to a measured train track T as shown in Figure 4. Here the measure assigned to the branch transverse to e_i is $\mathcal{L}(e_i)$ and the measure on the other branches is determined by the *switch condition* [19, Page 11]. Furthermore, if p is a path from $\mathcal{T}$ to $h(\mathcal{T})$ in G then this descends to a quasi-path in $\mathbb{G}$ and so $d(T, \widehat{h}(T)) \leq 2\ell(p)$.

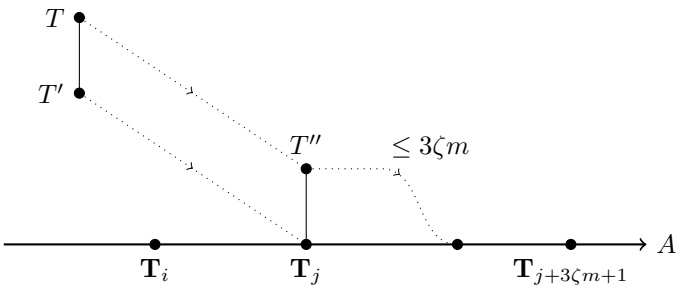

Figure 3: Getting to the axis A.

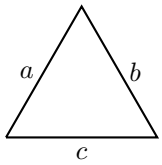

(a) When $\mathcal{L}(a) = \mathcal{L}(b) = \mathcal{L}(c) = 0$.

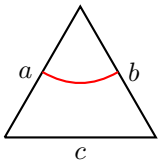

(b) When $\mathcal{L}(a) = \mathcal{L}(b)$ and $\mathcal{L}(c) = 0$.

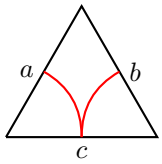

(c) When $\mathcal{L}(a) + \mathcal{L}(b) = \mathcal{L}(c)$.

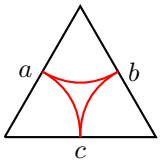

(d) Otherwise.

Figure 4: A train track coming from a triangulation.

3 Certification

We can now state the main algorithm for certifying that a mapping class is pseudo-Anosov. Again, we use a path p from $\mathcal{T}$ to $h(\mathcal{T})$ to represent h. We use a *certificate* consisting of decimals $x_1, \ldots, x_\zeta$ and polynomials $f_1, \ldots, f_\zeta \in \mathbb{Z}[x]$. These decimals represent approximations of the measure assigned to each edge of $\mathcal{T}$ by the stable lamination of h while these polynomials are their minimal polynomials.

Given a path p from $\mathcal{T}$ to $h(\mathcal{T})$ and a certificate $x_1, \ldots, x_\zeta, f_1, \ldots, f_\zeta$, let:

- $\{A_i\}$ and $\{B_i\}$ be the matrices of Lemma 1.6,
- $t := 24\zeta K \cdot \ell(p)^2$,
- $h_0 := \zeta^4(\ell(p) + 6)$,
- $h_1 := h_0 + 2\zeta$,
- $p_1 := 2\zeta^2(2h_1 + t + 3)$, and
- $d_1 := p_1 + t + \zeta h_1 + 2$.

In each of the following stages, all calculations are done to d_1 decimal places and all comparisons are done by *only* comparing the first p_1 decimal places.

1. Check that each f_i has degree at most ζ and the log of the absolute value of each coefficient is at most h_0.
2. Check that each $0 \leq x_i \leq 1$.
3. Check that each pair $f_i(x_i \pm 10^{-d_1})$ have different signs.
4. For each face of $\mathcal{T}$ with edges a, b, c check that $x_a + x_b \geq x_c$.
5. For each vertex $v \in V(\mathcal{T})$ check that there is an incident face of $\mathcal{T}$ with edges a, b, c such that $v \subseteq a \cap b$ and $x_a + x_b = x_c$.
6. Check that $\sum x_i = 1$.
7. Find B_i such that $B_i \cdot (x_1 \ \cdots \ x_\zeta)^T \geq 0$ and compute
$$(y_1 \ \cdots \ y_\zeta)^T := A_i \cdot (x_1 \ \cdots \ x_\zeta)^T \quad \text{and} \quad y := \sum y_i.$$
8. Check that each $y_i = yx_i$.
9. Check that $y > 1$.
10. Check that $s^t(T')$ is filling where T' is the measured train track corresponding to $x_1, \ldots, x_\zeta$ (see Section 2).

We say that a certificate is *accepted* by the main algorithm if every check passes and is *rejected* otherwise. We will show that:

- for each $h \in X^*$, there is a certificate that the main algorithm accepts if and only if h is pseudo-Anosov. (Theorem 3.18), and
- the main algorithm will accept or reject a certificate of h in $O(\ell(p)^4)$ time (Theorem 3.19).

As we may choose p such that $\ell(p) \in O(\ell(h))$, these two results suffice to show that the pseudo-Anosov problem is in **NP**.

3.1 Algebraic numbers

In order to prove the correctness of the main algorithm we first recall some properties of algebraic numbers.

Definition 3.1. If $f(x) = \sum a_i x^i \in \mathbb{Z}[x]$ is a polynomial then we define its *height* to be
$$\mathrm{hgt}(f) := \log(\max(|a_i|)).$$

Definition 3.2. If α is an algebraic number then we define its *height* to be $\mathrm{hgt}(\alpha) := \mathrm{hgt}(\mu_\alpha)$ where $\mu_\alpha \in \mathbb{Z}[x]$ is its minimal integral polynomial.

Fact 3.3. If $\alpha, \beta \in \overline{\mathbb{Q}}$ are algebraic numbers then:

- $\mathrm{dg}(\alpha \pm \beta) \leq \mathrm{dg}(\alpha) + \mathrm{dg}(\beta)$, where $\mathrm{dg}(\alpha) := \log(\deg(\alpha))$,
- $\mathrm{hgt}(\alpha \pm \beta) \leq \mathrm{hgt}(\alpha) + \mathrm{hgt}(\beta) + 1$ [23, Property 3.3],
- $\mathrm{hgt}(\alpha\beta) \leq \mathrm{hgt}(\alpha) + \mathrm{hgt}(\beta)$ [23, Property 3.3],
- $\mathrm{hgt}(\alpha^{-1}) = \mathrm{hgt}(\alpha)$, and
- if α is a root of $f \in \mathbb{Z}[x]$ then $\mathrm{hgt}(\alpha) \leq \mathrm{hgt}(f) + 2\deg(f)$ [3, Corollary 10.12].

Most importantly, algebraic numbers of bounded degree and height are bounded away from zero.

Lemma 3.4 ([8, Lemma 10.3])**.** *If α is an algebraic number then $\alpha = 0$ if and only if the integer part of α is 0 and at least the first $\mathrm{hgt}(\alpha) + \mathrm{dg}(\alpha)$ decimal places of α are 0.* □

Remark 3.5. In certain circumstances the inequalities of Fact 3.3 may be strengthened. For example, if $\alpha_1, \ldots, \alpha_k$ are integers then

$$\mathrm{hgt}\left(\sum \alpha_i\right) \leq \max(\mathrm{hgt}(\alpha_i)) + \log(k) \ll \sum \mathrm{hgt}(\alpha_i) + (k-1).$$

3.2 Acceptance implies pseudo-Anosov

We can now prove the correctness of the main algorithm. Namely that h is pseudo-Anosov if and only if there is a certificate that the main algorithm accepts. We begin with the converse direction and proceed by showing that:

- Stages 1–3 show that each x_i is close to an algebraic number (Proposition 3.6).
- Stages 4–5 show that these algebraic numbers correspond to a measured lamination (Proposition 3.7).
- Stage 6 shows that this lamination is *unitary (with respect to $\mathcal{T}$)*, that is, $||\mathcal{T}(\mathcal{L})|| = 1$, (Proposition 3.8).
- Stage 8 shows that this lamination is projectively invariant under h (Proposition 3.9).
- Stage 9 shows that this lamination is stable (Proposition 3.10).
- Stage 10 shows that this lamination is filling (Proposition 3.11).

Proposition 3.6. *If Stages 1–3 of the main algorithm complete then each x_i lies within 10^{-d_1} of a unique algebraic number v_i of degree at most ζ and height at most h_1.*

Proof. By the intermediate value theorem, Stage 3 shows that f_i must have a root v_i in $[x_i - 10^{-d_1}, x_i + 10^{-d_1}]$. It follows from Fact 3.3 that v_i is an algebraic number of degree at most $\deg(f_i) \leq \zeta$ and height at most $\mathrm{hgt}(f_i) + 2\zeta \leq h_0 + 2\zeta = h_1$. Finally, any two distinct algebraic number of degree at most ζ and height at most h_1 must be separated by at least $10^{-(2h_1+2\zeta+1)} > 10^{-(d_1-1)}$ and so v_i is unique. □

From now on fix $v_1, \ldots, v_\zeta$ to be these algebraic numbers.

Proposition 3.7. *Stages 4–5 of the main algorithm complete if and only if $v_1, \ldots, v_\zeta$ corresponds to a measured lamination $\mathcal{L} \in \mathcal{ML}(S)$, that is,*

$$\mathcal{T}(\mathcal{L}) = (v_1, \ldots, v_\zeta).$$

Proof. If $x_a + x_b > x_c$ to p_1 decimal places then

$$v_a + v_b - v_c > 10^{-p_1} - 10^{-d_1} - 10^{-d_1} \geq 10^{-(p_1-1)}.$$

However

$$\operatorname{hgt}(v_a + v_b - v_c) + \operatorname{dg}(v_a + v_b - v_c) \leq 3h_1 + 3\zeta \leq p_1 - 1$$

and so by Lemma 3.4 we have that $v_a + v_b > v_c$. By the same argument if $x_a + x_b = x_c$ to p_1 decimal places then $v_a + v_b = v_c$. Hence $v_1, \ldots, v_\zeta$ corresponds to a measured lamination.

Conversely, suppose that $v_1, \ldots, v_\zeta$ corresponds to a measured lamination. If $v_a + v_b > v_c$ then $v_a + v_b - v_c > 10^{-(3h_1+3\zeta)}$ and so

$$x_a + x_b - x_c > 10^{-(3h_1+3\zeta)} - 10^{-(d_1-1)} > 10^{-p_1}.$$

Hence $x_a + x_b > x_c$ to at least p_1 decimal places. Similarly if $v_a + v_b = v_c$ then $x_a + x_b = x_c$ to at least p_1 decimal places and so Stages 4–5 of the main algorithm will complete. □

From now on fix $\mathcal{L}$ to be this measured lamination.

Proposition 3.8. *Stage 6 of the main algorithm completes if and only if $\mathcal{L}$ is unitary (with respect to $\mathcal{T}$).*

Proof. If $\sum x_i = 1$ to p_1 decimal places then $\sum v_i - 1 = 0$ to at least $p_1 - 1$ decimal places. However

$$\operatorname{hgt}\left(\sum v_i - 1\right) + \operatorname{dg}\left(\sum v_i - 1\right) \leq \zeta h_1 + \zeta^2 \leq p_1 - 1$$

and so again by Lemma 3.4 we have that $\sum v_i = 1$. Hence $\mathcal{L}$ is unitary.

Conversely, if $\mathcal{L}$ is unitary then $\sum v_i - 1 = 0$ and so $|\sum x_i - 1| \leq \zeta 10^{-d_1}$. Hence $\sum x_i = 1$ to at least p_1 decimal places and so Stage 6 of the main algorithm will complete. □

Now following Stage 7, let B_i be such that $B_i \cdot (x_1 \ \cdots \ x_\zeta)^T \geq 0$ and fix

$$(y_1 \ \cdots \ y_\zeta)^T := A_i \cdot (x_1 \ \cdots \ x_\zeta)^T \quad \text{and} \quad y := \sum y_i.$$

As v_i and x_i are so close and comparisons are only done to p_1 decimal places we also have that $B_i \cdot (v_1 \ \cdots \ v_\zeta)^T \geq 0$. Hence we fix

$$(w_1 \ \cdots \ w_\zeta)^T := A_i \cdot (v_1 \ \cdots \ v_\zeta)^T \quad \text{and} \quad \lambda := \sum w_i.$$

Each A_i is $\ell(p)$–bounded and $\ell(p) \leq h_1$. Therefore each w_i has height at most $h_2 := \zeta(2h_1 + 1)$ and agrees with y_i to at least $d_2 := d_1 - \zeta h_1$ decimal places. Similarly, λ has height at most $h_3 := \zeta(h_2 + 1)$ and agrees with y to at least $d_3 := d_2 - \zeta$ decimal places.

Proposition 3.9. *Stage 8 of the main algorithm completes if and only if $\mathcal{L}$ is projectively invariant under h, that is, $h(\mathcal{L}) = \lambda \cdot \mathcal{L}$.*

Proof. Firstly note that $\lambda, y < 10^{h_3+\zeta}$ and so yx_i and λv_i agree to at least $d_3 - \zeta - h_3$ decimal places. Therefore, if $y_i - yx_i = 0$ to p_1 decimal places then $w_i - \lambda v_i = 0$ to at least $p_1 - 1$ places. However

$$\operatorname{hgt}(w_i - \lambda v_i) + \operatorname{dg}(w_i - \lambda v_i) \leq h_1 + h_2 + h_3 + 2\zeta^2 \leq p_1 - 1$$

and so again by Lemma 3.4 we have that $w_i = \lambda v_i$. Hence $\mathcal{L}$ is projectively invariant under h.

Conversely, if $\mathcal{L}$ is projectively invariant under h then $w_i - \lambda v_i = 0$ and so

$$|y_i - yx_i| \leq 10^{-d_2} \leq 10^{-p_1} + 10^{-(d_3-\zeta-h_3)}.$$

Hence $y_i = yx_i$ to at least p_1 decimal places and so Stage 8 of the main algorithm will complete. □

Note that completing this stage shows that $(v_1 \ \cdots \ v_\zeta)^T$ is an eigenvector of A_i. Hence each v_i lies in $\mathbb{Q}(\lambda)$ and so any linear combination of them is also an algebraic number of degree at most ζ.

Proposition 3.10. *Stage 9 of the main algorithm completes if and only if $\mathcal{L}$ is stable, that is, $\lambda > 1$.*

Proof. If $y - 1 > 0$ to p_1 decimal places then $\lambda - 1 > 0$ to at least $p_1 - 1$ decimal places. However

$$\operatorname{hgt}(\lambda - 1) + \operatorname{dg}(\lambda - 1) \leq h_3 + \zeta \leq p_1 - 1.$$

and so again by Lemma 3.4 we have that $\lambda > 1$. Hence $\mathcal{L}$ is stable.

Conversely if $\mathcal{L}$ is stable then $\lambda > 1 + 10^{-(h_3+\zeta)}$. As λ and y agree to at least d_3 decimal places we have that

$$y > 1 + 10^{-(h_3+\zeta)} - 10^{-d_3} > 1 + 10^{-p_1}.$$

Hence $y > 1$ to at least p_1 decimal places and so Stage 9 completes. □

Following Section 2, let T be the measured train track obtained from $\mathcal{T}$ using $v_1, \ldots, v_\zeta$ and let $T_i := s^i(T) = (\tau_i, \mu_i)$. Similarly, let T' be the measured train track obtained from $\mathcal{T}$ using $x_1, \ldots, x_\zeta$ instead and let $T'_i := s^i(T') = (\tau'_i, \mu'_i)$. We let v_i^k denote the weights on the branches of T_k and x_i^k denote the weights on the branches of T'_k.

Proposition 3.11. *Stage 10 of the main algorithm completes if and only if $\mathcal{L}$ is filling.*

Proof. We will show that T'_t is filling if and only if T_t is. The result then follows directly from Corollary 2.13.

We begin by claiming that for $1 \leq k \leq t$:

$$\tau_k = \tau'_k \quad \text{and} \quad |v_i^k - x_i^k| \leq 10^{-(d_1-k-1)}.$$

To see this first note that v_i^k is an algebraic number of degree at most ζ and $\operatorname{hgt}(v_i^k) \leq 3\zeta(k + 2h_1 + 2)$ and that $|v_i^0 - x_i^0| \leq 10^{-(d_1-1)}$. Now suppose that

$\tau_k = \tau'_k$ and $|v_i^k - x_i^k| \leq 10^{-(d_1-k-1)}$ for some $1 \leq k < t$. Then by Lemma 3.4, $v_i^k \geq v_j^k$ if and only if $x_i^k \geq x_j^k$ to p_1 decimal places. Therefore the x_i^k–maximal branches are the v_i^k–maximal branches and so $\tau_{k+1} = \tau'_{k+1}$. Furthermore

$$|v_i^{k+1} - x_i^{k+1}| \leq |v_i^k - x_i^k| + |v_j^k - x_j^k| \leq 10^{-(d_1-k-2)}$$

and so the claim holds by induction on k.

Finally, again by Lemma 3.4, we have that $v_i^t > 0$ if and only if $x_i^t > 0$ to p_1 decimal places. Hence T_t is filling if and only if T'_t is and so Stage 10 completes if and only if $\mathcal{L}$ is filling. □

Combining these propositions we obtain:

Corollary 3.12. *Suppose that $h \in \mathrm{Mod}^+(S)$ is a mapping class and p is a path from $\mathcal{T}$ to $h(\mathcal{T})$. If there is a certificate that the main algorithm accepts then h is pseudo-Anosov.* □

3.3 Pseudo-Anosovs have acceptable certificates

Finally, we show the converse to Corollary 3.12. To do this we first require some additional bounds on the heights of certain algebraic numbers.

Definition 3.13. Suppose that $\alpha \in \overline{\mathbb{Q}}$ is an algebraic number. A matrix M is *α–shifted* if its entries are of the form $a_{ij} = b_{ij} + c_{ij}\alpha$, where b_{ij} and c_{ij} are integers. We say that such a matrix is *k–bounded* if each b_{ij} and c_{ij} is.

Proposition 3.14. *If M is a k–bounded, $m \times m$, α–shifted matrix then*

$$\mathrm{hgt}(\det(M)) \leq m^2(k + \log(m) + \mathrm{hgt}(\alpha)).$$

Proof. First note that we may expand $(b_1 + c_1\alpha) \cdots (b_m + c_m\alpha)$ as a polynomial in α to obtain $d_0 + \cdots + d_m\alpha^m$. It then follows from Fact 3.3 that $\mathrm{hgt}(d_i) \leq mk$ as b_i and c_i are k–bounded integers .

Now consider the following expansion of $\det(M)$:

$$\det(M) = \sum_{\sigma \in \mathrm{Sym}(m)} \mathrm{sign}(\sigma) \prod_{i=1}^{m} \left(b_{i\sigma(i)} + c_{i\sigma(i)}\alpha\right) = e_0 + \cdots + e_n\alpha^m.$$

By applying the previous bound to the coefficients of $\prod_{i=1}^{m} \left(b_{i\sigma(i)} + c_{i\sigma(i)}\alpha\right)$ we have that

$$\mathrm{hgt}(e_i) \leq mk + m\log(m).$$

Therefore:

$$\begin{aligned}
\mathrm{hgt}(\det(M)) &\leq \mathrm{hgt}\left(\sum e_i\alpha^i\right) \\
&\leq \sum \mathrm{hgt}(e_i\alpha^i) + m\log(2) \\
&\leq \sum \mathrm{hgt}(e_i) + \frac{1}{2}m^2\,\mathrm{hgt}(\alpha) + m\log(2) \\
&\leq m^2k + m^2\log(m) + \frac{1}{2}m^2\,\mathrm{hgt}(\alpha) + m\log(2) \\
&\leq m^2(k + \log(m) + \mathrm{hgt}(\alpha))
\end{aligned}$$

Thus showing the required bound. □

Lemma 3.15. *Suppose that M is a k–bounded, $m \times m$, α–shifted matrix and that $\det(M) \neq 0$. If $v = (\alpha_1 \ \cdots \ \alpha_m)^T$ is a vector of algebraic numbers such that $M \cdot v$ is a k–bounded vector of integers then*

$$\mathrm{hgt}(\alpha_i) \leq 2m^2(k + \log(m) + \mathrm{hgt}(\alpha))$$

Proof. This bound follows from using Cramer's rule to determine α_i and Proposition 3.14. □

Lemma 3.16. *Suppose that M is a k–bounded $m \times m$ integer matrix. If α is an eigenvalue of M then $\mathrm{hgt}(\alpha) \leq mk + m\log(m) + 2m$.*

Proof. Again, following the proof of Proposition 3.14, we have that if $\chi_M(x) = e_0 + \cdots + e_n x^m$ is the characteristic polynomial of M then $\mathrm{hgt}(e_i) \leq mk + m\log(m)$. As α is a root of this it follows from Fact 3.3 that

$$\mathrm{hgt}(\alpha) \leq mk + m\log(m) + 2m. \qquad \square$$

We now have the tools to prove existence of a certificate that the main algorithm will accept for pseudo-Anosovs.

Theorem 3.17. *Suppose that $h \in \mathrm{Mod}^+(S)$ is a mapping class and p is a path from $\mathcal{T}$ to $h(\mathcal{T})$. If h is pseudo-Anosov then there is a certificate that the main algorithm will accept.*

Proof. Let $\mathcal{L} = \mathcal{L}^+(h)$ be the stable lamination of h scaled such that $||\mathcal{T}(\mathcal{L})|| = 1$. Let $v_i := \mathcal{L}(e_i)$, where e_i are the edges of $\mathcal{T}$, and $\mathbf{v} := (v_1 \ \cdots \ v_\zeta)^T$. Let $\{A_i\}$ and $\{B_i\}$ be the matrices of Lemma 1.6. There is B_i such that $B_i \cdot \mathbf{v} \geq 0$ and so $\mathbf{v}$ is an eigenvector of A_i. Therefore, as A_i is $\ell(p)$–bounded, each v_i is an algebraic number of degree at most ζ and by Lemma 3.15 and Lemma 3.16 we have that

$$\begin{aligned}
\mathrm{hgt}(v_i) &\leq 2\zeta^2(\ell(p) + \log(\zeta) + \zeta\ell(p) + \zeta\log(\zeta) + 2\zeta) \\
&\leq 2\zeta^2((\zeta+1)\ell(p) + (\zeta+1)\log(\zeta) + 2\zeta) \\
&\leq \zeta^4(\ell(p) + 6) \\
&= h_0.
\end{aligned}$$

Now let x_i be a decimal approximation of v_i, correct to d_1 decimal places, and let f_i be the minimal integral polynomial of v_i. We claim that the certificate

$$x_1, \ldots, x_\zeta, f_1, \ldots, f_\zeta$$

is accepted by the main algorithm. To see this consider the first three stages of the main algorithm:

1. By definition each f_i has degree at most ζ and height at most h_0. Hence this stage will pass.
2. As $0 \leq v_i \leq 1$ and $|x_i - v_i| \leq 10^{-d_1}$, we have that $0 \leq x_i \leq 1$ to at least p_1 decimal places. Hence this stage will pass.
3. By Fact 3.3, two distinct roots of f_i must be separated by at least $10^{-(2h_1+\zeta)}$ and as f_i is minimal it has no repeated roots. Hence, v_i is the unique root of f_i in $[x_i - 10^{-d_1}, x_i + 10^{-d_1}]$ and so $f_i(x_i \pm 10^{-d_1})$ must have different signs. Hence this stage will pass and the algebraic numbers found will be $v_1, \ldots, v_\zeta$.

Finally, note that as $v_1, \ldots, v_\zeta$ represents $\mathcal{L}$ which is unitary (with respect to $\mathcal{T}$), projectively invariant, stable and filling Stages 6–10 must complete by Proposition 3.8, Proposition 3.9, Proposition 3.10 and Proposition 3.11. Proving the claim that the main algorithm will accept this certificate. □

Together with Corollary 3.12 this shows:

Theorem 3.18. *Suppose that $h \in \mathrm{Mod}^+(S)$ is a mapping class and p is a path from $\mathcal{T}$ to $h(\mathcal{T})$. There is a certificate that the main algorithm will accept if and only if h is pseudo-Anosov.* □

3.4 Analysis

Finally, we consider the running time of the main algorithm.

Theorem 3.19. *The main algorithm will accept or reject a certificate of h in $O(\ell(p)^4)$ time.*

Proof. We analyse each of the stages of the main algorithm in turn. Recall that $h_0, h_1 \in O(\ell(p))$ and $t, d_1, p \in O(\ell(p)^2)$.

1. This can be done in $O(h_0) = O(\ell(p))$ operations.
2. This can be done in $O(p_1) = O(\ell(p)^2)$ operations.
3. By expanding $f_i(x) = \sum a_j x^j$ as $a_0 + x(\cdots(a_{\zeta-2} + x(a_{\zeta-1} + a_\zeta x))\cdots)$ using Horner's rule [12, Section 4.6.4] it can be seen that each $f_i(x_i \pm 10^{-d_1})$ can be computed in $O(d_1(d_1 + \zeta h_0)) = O(\ell(p)^4)$. Hence this stage can be done in $O(\ell(p)^4)$ operations.
4. There are $O(1)$ inequalities to check each of which requires $O(d_1 + p_1) = O(\ell(p)^2)$ operations. Hence this can be done in $O(\ell(p)^2)$ operations.
5. By the same argument this also requires $O(\ell(p)^2)$ operations.
6. This can be done in $O(d_1 + p_1) = O(\ell(p)^2)$ operations.
7. Following the ideas of Remark 1.7, as each x_i has d_1 digits, we can find A_i and B_i and compute $y_1, \ldots, y_\zeta$ in $d_1\ell(p) + \ell(p)^2 \in O(\ell(p)^3)$ operations. Each y_i has at most $d_1 + \ell(p) = O(\ell(p)^2)$ digits. Hence, y can be computed in $O(\ell(p)^2)$ further operations and has $O(\ell(p)^2)$ digits. Therefore this entire stage can be done in $O(\ell(p)^3)$ operations.
8. As x_i, y_i and y each have $O(\ell(p)^2)$ digits this stage can be done in $O(\ell(p)^4 + d_1 + p_1) = O(\ell(p)^4)$ operations.
9. As y has $O(\ell(p)^2)$ digits, this can be done in $O(\ell(p)^2 + p_1) = O(\ell(p)^2)$ operations.
10. Constructing T' takes at most $O(d_1) = O(\ell(p)^2)$ operations. For each maximal splitting it takes $O(p_1)$ operations to find the maximal weight branches and then $O(d_1)$ further operations to perform the splitting. Therefore we can construct $T'_t := s^t(T')$ in $O(t(p_1 + d_1)) = O(\ell(p)^4)$ operations. Finally, checking whether T'_t is filling can be done in $O(p_1) = O(\ell(p)^2)$ operations. Hence this whole stage can be done in $O(\ell(p)^4)$ time.

Therefore the main algorithm will terminate in $O(\ell(p)^4)$ time. □

Again, we may choose p such that $\ell(p) \in O(\ell(h))$. Hence we deduce that:

Corollary 3.20. *Fix S, a marked surface, and X, a finite generating set of* $\mathrm{Mod}^+(S)$. *Deciding whether a word $h \in X^*$ correspond to a pseudo-Anosov mapping classes is a problem in **NP**.* □

As usual, this immediately gives an exponential time algorithm to determine whether a word is pseudo-Anosov.

4 The conjugacy problem

In this section we apply the results developed in Section 2 to the conjugacy problem for mapping class groups. We use these to show that that the conjugacy problem is in **co-NP** (Corollary 4.15). That is, if two words do not represent conjugate mapping classes then there is a short proof of this.

To do this we work with a path p from $\mathcal{T}$ to $h(\mathcal{T})$ and deal with three cases depending on the Nielsen–Thurston type of h. In each case we will describe a total conjugacy invariant and show that this can be constructed in $O(\mathrm{poly}(\ell(p)))$ time when given a suitable certificate. Thus, if q is a second path from $\mathcal{T}'$ to $g(\mathcal{T}')$ then these invariants are small enough they can be compared in $O(\mathrm{poly}(\ell(p) + \ell(q)))$ operations. Hence we can determine whether g and h are conjugate in polynomial time.

Again we may choose paths p and q such that $\ell(p) \in O(\ell(h))$ and $\ell(q) \in O(\ell(g))$. Hence this shows that there are certificates that allow us to determine whether g and h are conjugate in $O(\mathrm{poly}(\ell(g) + \ell(h)))$ operations.

The reducible case relies on a solution to a stronger version of the conjugacy problem for the other two cases. Specifically, a solution to the *permutation conjugation problem*:

Problem 4.1 (The permutation conjugacy problem)**.** Suppose that $g, h \in \mathrm{Mod}^+(S)$ are mapping classes. Given paths p from $\mathcal{T}$ to $h(\mathcal{T})$ and q from $\mathcal{T}$ to $g(\mathcal{T}')$ and a map $\pi\colon V \subseteq V(\mathcal{T}) \to V(\mathcal{T}')$, decide whether g and h are π–*conjugate*. That is, decide whether there is there is a mapping class $f \in \mathrm{Mod}^+(S)$ such that $fgf^{-1} \equiv h$ and $f|_V = \pi$.

4.1 Periodic mapping classes

We begin with the case in which h is periodic. Here we use properties from its quotient orbifold as a total conjugacy invariant.

As h is periodic, by the Nielsen realization theorem [10, Theorem 7.2] there is a homeomorphism $\phi \in h$ such that $\mathrm{ord}(\phi) = \mathrm{ord}(h)$. We may use this to define the *quotient orbifold* $\mathcal{O} := S \,/\, \phi$ which, up to homeomorphism, is independent of the particular choice of ϕ. There are two key properties that we will need to determine $\mathcal{O}$.

Firstly, we say that the *order* of a point $x \in S$ is

$$\mathrm{ord}(x) := \begin{cases} \min\{k > 0 : \phi^k(x) = x\} & \text{if } x \text{ is not a marked point,} \\ -\min\{k > 0 : \phi^k(x) = x\} & \text{otherwise.} \end{cases}$$

We note that all but finitely many of the points of S are *regular*, that is, have order $\mathrm{ord}(\phi)$. In fact the number of irregular points is bounded only in terms of the topology of S.

Secondly, even though $\phi^{|\mathrm{ord}(x)|}(x) = x$ this map may not act like the identity near x; it may rotate the tangent plane at x by $2\pi r(x)$. We refer to the rational number $r(x)$ as the *rotation number* of x. We note that the denominator of $r(x)$ is bounded above by $\mathrm{ord}(\phi) \leq 8\,\mathrm{g}(S) + 4\,\mathrm{n}(S) - 2$ [10, Theorem 7.5].

Let $N(o, r)$ denote the number of irregular points of S with order o and rotation number r. Nielsen showed that N determines $\mathcal{O}$ and so is a total conjugacy invariant of h [18, Theorem 9].

To compute N we consider a *multiarc*, that is, is the isotopy class of the image of a smooth proper embedding of a finite number of copies of $[0, 1]$ (whose endpoints are sent to marked points) into S. Again we represent a multiarc via its intersection numbers with the edges of a triangulation. However, as a non-trivial multiarc can have zero intersection with all edges, we must first make a slight modification to the standard definition of intersection number:

Definition 4.2. If α is a multiarc which contains k copies of the edge e of $\mathcal{T}$ then their *intersection number* is defined to be $\iota(\alpha, e) := -k$.

Having made this change we can now identify a multiarc with its *normal coordinate*, its vector of intersection numbers with the edges of $\mathcal{T}$. These coordinates allow us to restate many of the results of Section 1.2 for multiarcs.

Firstly, as in Lemma 1.4, there are $O(1)$–bounded $\zeta \times 3\zeta$ matrices F_i' such that a vector v corresponds to a multiarc if and only if $v \neq 0$ and

$$F_i' \cdot v \geq_2 0$$

for some i. Secondly, analysis of the 30 cases that can occur within a pair of triangles shows how these coordinates change under an edge flip.

Lemma 4.3. *Suppose that α is a multiarc and e is a flippable edge of a triangulation $\mathcal{T}$ as shown in Figure 1 then*

$$\iota(\alpha, e') = \begin{cases} \widehat{a} + \widehat{b} - \bar{e} & \text{if } \bar{e} \geq \widehat{a} + \widehat{b} \text{ and } \widehat{a} \geq \widehat{d} \text{ and } \widehat{b} \geq \widehat{c}, \\ \widehat{c} + \widehat{d} - \bar{e} & \text{if } \bar{e} \geq \widehat{c} + \widehat{d} \text{ and } \widehat{d} \geq \widehat{a} \text{ and } \widehat{c} \geq \widehat{b}, \\ \widehat{a} + \widehat{d} - \bar{e} & \text{if } \bar{e} \leq 0 \text{ and } \widehat{a} \geq \widehat{b} \text{ and } \widehat{d} \geq \widehat{c}, \\ \widehat{b} + \widehat{c} - \bar{e} & \text{if } \bar{e} \leq 0 \text{ and } \widehat{b} \geq \widehat{a} \text{ and } \widehat{c} \geq \widehat{d}, \\ \widehat{a} + \widehat{d} - 2\bar{e} & \text{if } \bar{e} \geq 0 \text{ and } \widehat{a} \geq \widehat{b} + \bar{e} \text{ and } \widehat{d} \geq \widehat{c} + \bar{e}, \\ \widehat{b} + \widehat{c} - 2\bar{e} & \text{if } \bar{e} \geq 0 \text{ and } \widehat{b} \geq \widehat{a} + \bar{e} \text{ and } \widehat{c} \geq \widehat{d} + \bar{e}, \\ \frac{1}{2}(\widehat{a} + \widehat{b} - \bar{e}) & \text{if } \widehat{a} + \widehat{b} \geq \bar{e} \text{ and } \widehat{b} + \bar{e} \geq 2\widehat{c} + \widehat{a} \text{ and } \widehat{a} + \bar{e} \geq 2\widehat{d} + \widehat{b}, \\ \frac{1}{2}(\widehat{c} + \widehat{d} - \bar{e}) & \text{if } \widehat{c} + \widehat{d} \geq \bar{e} \text{ and } \widehat{d} + \bar{e} \geq 2\widehat{a} + \widehat{c} \text{ and } \widehat{c} + \bar{e} \geq 2\widehat{b} + \widehat{d}, \\ \max(\widehat{a} + \widehat{c}, \widehat{b} + \widehat{d}) - \bar{e} & \text{otherwise} \end{cases}$$

where $\widehat{x} := \max(\iota(\alpha, x), 0)$ and $\bar{e} := \iota(\alpha, e)$. □

Now suppose that α is a multiarc in $\mathcal{O}$ which *fills*, that is, such that each component of $\mathcal{O} - \alpha$ is a disk containing at most one irregular point in its interior. For example, see Figure 5a. By lifting α back to S we see that there is an h–invariant filling multiarc β on S. Now β decomposes S into polygonal

pieces, each of which contains at most one irregular point. For example, see Figure 5b. These polygons are permuted by the action of h and it is from this permutation that we can immediately determine the orders and rotation numbers of the irregular points of S and so compute N.

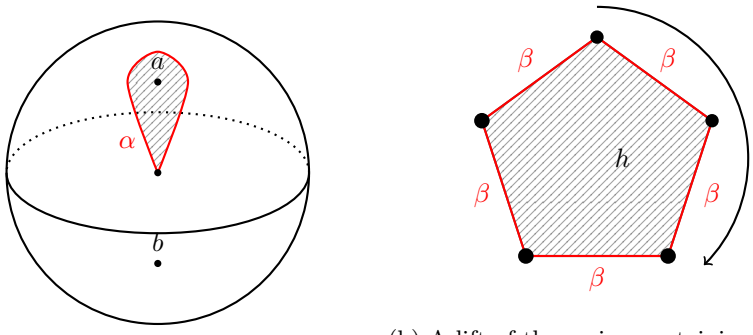

(a) A filling multiarc α on $\mathcal{O}$.

(b) A lift of the region containing a to S when $\text{ord}(a) = 5$.

Figure 5: An invariant multiarc β obtained by lifting a multiarc α.

Thus, given a path p from $\mathcal{T}$ to $h(\mathcal{T})$ and vector $v := \mathcal{T}(\beta)$ we:

1. Check that v corresponds to a multiarc β by checking that $F'_i \cdot v \geq_2 0$ for some i.
2. Check that β is h–invariant by computing and checking that $\mathcal{T}(h(\beta)) = v$.
3. Compute the normal coordinate of each component β_i of β and check that each occurs with multiplicity one.
4. Compute the components β_i that are adjacent to each region of $S - \beta$, together with their cyclic ordering.
5. Compute the permutation of β_i induced by h and so compute N.

To analyse the number of operations required by this procedure, suppose that the given vector v is k–bounded. Then checking that $F'_i \cdot v \geq_2 0$ for some i can be done in $O(\text{poly}(k))$ operations as the F'_i matrices are $O(1)$–bounded. By repeating the idea of Remark 1.7, we can compute $\mathcal{T}(h(\beta))$ and so complete Stage 2 in $O(\text{poly}(k)\,\text{poly}(\ell(p)))$ operations.

To complete Stage 3 and Stage 4 we can use the algorithm of Agol, Hass and Thurston [2] or one of its variants [9, Section 6] [20, Theorem 1]. These allow us to extract the edge vectors of the components of β in $O(\text{poly}(k))$ operations. Finally by computing $h(\beta_i)$ using the idea of Remark 1.7, this permutation can be computed in $O(\text{poly}(k)\,\text{poly}(\ell(p)))$ operations and from this N can be computed in $O(\text{poly}(k)\,\text{poly}(\ell(p)))$ operations too.

Now if p is a path from $\mathcal{T}$ to $h(\mathcal{T})$ we can once again compute matrices A'_i and B'_i which describe the piecewise linear transformation between coordinates with respect to $\mathcal{T}$ and $h(\mathcal{T})$. Again, these matrices can be chosen such that:

- each A'_i and B'_i is $\ell(p)$–bounded,
- each B'_i has $O(\ell(p))$ rows,

- for each multiarc α we have that $B'_i \cdot \mathcal{T}(\alpha) \geq 0$ for some i, and
- for each multicurve α we have that

$$\mathcal{T}(h(\alpha)) = A'_i \cdot \mathcal{T}(\alpha) \quad \text{if and only if} \quad B'_i \cdot \mathcal{T}(\alpha) \geq 0.$$

We can use A'_i, B'_i and F'_i to repeat the argument of [4, Theorem 3.3]. From this we deduce that there is a filling multiarc β on S such that $h(\beta) = \beta$ and $\mathcal{T}(\beta)$ is $O(\ell(p))$–bounded. Hence there is a v such that we can compute N in $O(\text{poly}(\ell(p)))$ operations.

Corollary 4.4. *Deciding whether two periodic words $g, h \in X^*$ correspond to conjugate mapping classes is a problem in* ***co-NP****.* □

To deal with the permutation conjugacy problem we use the following lemma.

Lemma 4.5. *Periodic mapping classes h and g are π–conjugate if and only if:*

- *h and g are conjugate, and*
- *for each marked point $v \in V(S)$ we have that:*
  - *v and $\pi(v)$ have the same rotation number (with respect to h and g respectively), and*
  - *$\pi(h^k(v)) = g^k(\pi(v))$ for every $k \in \mathbb{Z}$*

  whenever these maps are defined.

Proof. The forward direction of this lemma holds trivially. For the reverse direction, we first note that without loss of generality we may assume that π is actually a permutation of the vertices of S. If it is not then we consider each of the possible extensions of π in turn.

As h and g are conjugate there is a mapping class f such that $h = f^{-1}gf$. If $f|_V \neq \pi$ then consider $\varphi\colon S/h \to S/g$, the homeomorphism induced by f. This homeomorphism respects the orders and rotation numbers of points and satisfies the *lifting criterion*: it maps the subgroup $\pi_1(S) \leq \pi_1(S/h)$ to the subgroup $\pi_1(S) \leq \pi_1(S/g)$. We can modify φ by precomposing it with another homeomorphism $\psi\colon S/h \to S/h$. If we take ψ to be a homeomorphism that swaps two of the marked points whose lifts have the same rotation number then ψ preserves the subgroup $\pi_1(S) \leq \pi_1(S/h)$. Hence $\varphi \circ \psi$ also satisfies the lifting criterion and so lifts to an alternate mapping class f' which, like f, conjugates g to h but whose action on V is permuted by ψ.

Therefore, as π sends marked points to marked points with the same rotation number and $\pi(h^k(v)) = g^k(\pi(v))$ for every $k \in \mathbb{Z}$, modifications of this form are sufficient to adjust f such that $f|_V = \pi$. Hence h and g are π–conjugate. □

The rotation numbers of the marked points can be determined in polynomial time from the polygonal decomposition of S given by the multiarc β. Thus this additional criterion can also be tested in polynomial time.

4.2 Pseudo-Anosov mapping classes

Agol showed that for pseudo-Anosov mapping classes the combinatorics of the periodic part of their maximal splitting sequence is a total conjugacy invariant [1, Section 7]. Thus, as we can construct these in polynomial time when given an acceptable certificate for p, we immediately obtain that:

Corollary 4.6. *Deciding whether two aperiodic irreducible words $g, h \in X^*$ correspond to conjugate mapping classes is a problem in* ***co-NP****.* □

We also note that in this case g and h are π–conjugate if and only if they are conjugate and there is a map between the combinatorics of the periodic part of their maximal splitting sequence which induces π on V. Again, given acceptable certificates for p and q, this can be done in $O(\text{poly}(\ell(p) + \ell(q)))$ operations.

4.3 Reducible mapping classes

To deal with the final case in which h is reducible (and aperiodic) we use its partition graph as a total conjugacy invariant. This involves decomposing the given mapping class along its canonical curve system and considering the maps induced on the invariant subsurfaces.

4.3.1 Crushing

In order to be able to construct the map induced on invariant subsurfaces computationally, we first recall the notion of *crushing* a surface along a curve from [4, Section 4].

Definition 4.7. We *crush* S along γ to obtain the (possibly disconnected) surface S_γ by removing an open regular neighbourhood of γ and then collapsing the new boundary components to additional marked points. Note that in order for this surface to be triangulable we must also discard any components that are twice marked spheres. See Figure 6 for example.

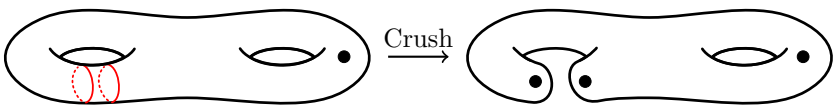

Figure 6: Crushing along a multicurve.

Now if $\mathcal{T}$ is a triangulation of S then we may track it as we crush S along a multicurve $\gamma \in \mathcal{C}(S)$. After collapsing any bigons that are created, this results in a triangulation $\mathcal{T}_\gamma$ of S_γ which also has ζ edges. Erickson and Nayyeri showed how to construct $\mathcal{T}_\gamma$ in polynomial time:

Theorem 4.8 ([9])**.** *There is an algorithm to compute $\mathcal{T}_\gamma$ which completes in $O(\text{poly}(k))$ operations when $\mathcal{T}(\gamma)$ is k–bounded.* □

Proposition 4.9. *If p is a path from $\mathcal{T}$ to $\mathcal{T}'$ then crushing each triangulation of p along γ, and possibly discarding any repeated triangulations, gives a path p_γ from $\mathcal{T}_\gamma$ to $\mathcal{T}'_\gamma$.*

Proof. The result clearly holds when p consists of a single reordering of the edges of $\mathcal{T}$. If p consists of a single flip then the combinatorics of $\mathcal{T}_\gamma$ and $\mathcal{T}'_\gamma$ agree away from the faces coming from the faces incident to the flipped edge. Thus $\mathcal{T}_\gamma$ and $\mathcal{T}'_\gamma$ share at least $\zeta - 1$ edges and so they are either equal or differ by a single flip. The result then follows for all paths by induction on $\ell(p)$. □

In fact when $\mathcal{T}'$ is obtained by flipping the edge e of $\mathcal{T}$, we have that $\mathcal{T}_\gamma$ and $\mathcal{T}'_\gamma$ are equal if and only if there is an arc of γ passing from one side of the square containing e to the opposite side. Following the notation of Figure 1, this occurs if and only if $\iota(\gamma, a) + \iota(\gamma, c) \neq \iota(\gamma, b) + \iota(\gamma, d)$.

We note that by construction $\ell(p_\gamma) \leq \ell(p)$.

Corollary 4.10. *Suppose that p is a path from $\mathcal{T}$ to $\mathcal{T}'$ and $\gamma \in \mathcal{C}(S)$ is a k–bounded multicurve. If $\mathcal{T}(\gamma)$ is k–bounded then we can compute p_γ in $O(\ell(p)\,\mathrm{poly}(k))$ operations.* □

4.3.2 Partition graphs

Now let $\sigma(h) \neq \emptyset$ denote the *canonical curve system* of h [10, Page 373]. This is the intersection of all maximal h–invariant multicurves. We will assume that $\sigma(h)$ has components $\{\gamma_j\}$. Let $S_{\sigma(h)}$ be the surface obtained by crushing S along $\sigma(h)$ and for ease of notation let $\{S_i\}$ be its connected components. Additionally, let h_i be the mapping class induced on S_i by the first return map of h.

Definition 4.11 ([14, Theorem 2], [18])**.** The *partition graph* of h is the pair (H, ϕ) where:

- H is the finite graph with:
  - a vertex corresponding to each h_i, and
  - an edge corresponding to each γ_i, connecting between h_j and h_k when γ_i meets S_j and S_k.
- ϕ is the automorphism of H induced by h.

Most importantly, we can provide a polynomial-time verifiable certificate that a given graph is the partition graph of h. This certificate consists of two pieces of information. Firstly $\mathcal{T}(\sigma(h))$ which is $O(\ell(p))$–bounded [4, Corollary 4.9] [13, Theorem 1.1]. From this we can compute paths representing each h_i in $O(\mathrm{poly}(\ell(p)))$ operations by using the algorithm of Theorem 4.8. Secondly a certificate accepted by the main algorithm for each pseudo-Anosov h_i, this allows us to deduce that the given multicurve is h–maximal.

Such information also allows us to verify that the given multicurve is actually the canonical curve system. To do this we note that it is sufficient to check that for any γ_i, removing its orbit under h from $\sigma(h)$ does not result in a new periodic component. As there are at most $\zeta \in O(1)$ such orbits to check and again we can compute each $\mathcal{T}(\gamma_i)$ in $O(\mathrm{poly}(\ell(p)))$ operations [9, Section 6] [2] [20, Theorem 1], this can also be done in polynomial time.

4.3.3 Twist invariants

Associated to each component γ_i of the canonical curve system is its *twist invariant* d_i. This is a rational number describing the number of (fractional) Dehn twists performed about it [18]. We now explicitly describe how to compute these numbers given some additional multicurves.

Definition 4.12. A curve γ is *dual* to a component γ_i of $\sigma(h)$ if

$$\iota(\gamma, \gamma_i) = \begin{cases} 2 & \text{if } \gamma_i \text{ is separating} \\ 1 & \text{otherwise.} \end{cases}$$

Now suppose that γ is dual to γ_i and let $\delta := \partial N(\gamma_i \cup \gamma)$. Let $k := \operatorname{lcm}(1, 2, \ldots, 4\zeta)$ then

$$|d_i| = \frac{\iota(h^k(\gamma), \gamma) - \frac{1}{2}\,\iota(h^k(\gamma), \delta)}{k\,\iota(\gamma, \gamma_i)}.$$

This particular value of k was chosen such that h^k fixes all components of $\sigma(h)$ and all prongs of singularities of any measured lamination that is projectively invariant under h.

To compute the sign of d_i we check whether the number of intersections grows when an additional Dehn twist along γ_i is performed. That is, $d_i \geq 0$ if and only if

$$\iota(h^k(T_{\gamma_i}(\gamma)), \gamma) \geq \iota(h^k(\gamma), \gamma).$$

To compute these intersection numbers we use the algorithm of Schaefer, Sedgwick and Štefankovič [21, Section 5]. This computes the intersection number of two multicurves from their edge vectors and does so in polynomial time [21, Lemma 5.4].

Thus given the edge vectors of a γ and δ, we can use the algorithm of Schaefer, Sedgwick and Štefankovič to first compute $\iota(\gamma, \gamma_i)$ and so verify that γ is dual to γ_i. Furthermore we can crush S along δ and check that the component of S_δ containing γ_i and γ is either a four times marked sphere or a once-marked torus, depending on whether γ_i is separating or not. This verifies that $\delta = \partial N(\gamma_i \cup \gamma)$. Having established that γ and δ are the relevant curves, we can use them to compute d_i by the previous formulae.

Now as $\mathcal{T}(\gamma_i)$ is $O(\ell(p))$–bounded there is a dual curve γ such that $\mathcal{T}(\gamma)$ is also $O(\ell(p))$–bounded. For such a γ, the edge vector of the corresponding δ is also $O(\ell(p))$–bounded. Therefore these intersection number and S_δ can all be computed in $O(\operatorname{poly}(\ell(p)))$ operations by the algorithm of Schaefer, Sedgwick and Štefankovič and Theorem 4.8. Hence we can compute d_i in $O(\operatorname{poly}(\ell(p)))$ operations too.

4.3.4 Equivalence of partition graphs

Now suppose that (G, ϕ) and (H, ψ) are the partition graphs of g and h respectively. A graph isomorphism $\Phi\colon G \to H$ induces a bijection between the marked points of g_i coming from $\sigma(g)$ and those of $\Phi(g_i)$ coming from $\sigma(h)$. We denote this by $D_{g_i}\Phi$.

Theorem 4.13 ([18, Theorem 8.3] [17, Section 1.14])**.** *Suppose that (G, ϕ) and (H, ψ) are the partition graphs of g and h respectively. Then g and h are conjugate if and only if there is a graph isomorphism $\Phi\colon G \to H$ such that:*

1. *Φ conjugates ϕ to ψ,*
2. *g_i and $\Phi(g_i)$ are $D_{g_i}\Phi$–conjugate, and*
3. *γ_j and $\Phi(\gamma_j)$ have the same twist invariant.* □

Thus when g and h are not conjugate, for each graph isomorphism $\Phi\colon G \to H$ either:

1. $\Phi \circ \phi \neq \psi \circ \Phi$, which we can check in $\zeta \in O(1)$ operations,
2. there is a vertex $g_i \in G$ such that g_i and $\Phi(g_i)$ are not $D_{g_i}\Phi$–conjugate and so by Corollary 4.4 and Corollary 4.6 we can provide a polynomial-time checkable proof of this fact, or
3. there is an edge $\gamma_j \in G$ such that γ_j and $\Phi(\gamma_j)$ have different twist invariants and so by Lemma 4.3.3 we can provide a polynomial-time checkable proof of this fact.

By doing this for each of the at most $\zeta! \in O(1)$ such isomorphisms, we can provide a proof that g and h are not conjugate which can be verified in $O(\text{poly}(\ell(p)+\ell(q)))$ operations.

Again, applying this result to the standard path of a word we obtain that:

Corollary 4.14. *Deciding whether two aperiodic, reducible words $g, h \in X^*$ correspond to conjugate mapping classes is a problem in **co-NP**.* □

Together with Corollary 4.4 and Corollary 4.6 this shows that:

Corollary 4.15. *Fix S, a marked surface, and X, a finite generating set of $\text{Mod}^+(S)$. Deciding whether two words $g, h \in X^*$ correspond to conjugate mapping classes is a problem in **co-NP**.* □

Finally, we note that again this immediately gives an exponential time algorithm to determine whether two words represent conjugate mapping classes.

5 Implementation

The main algorithm has been implemented as part of `flipper` [5].

However the implementation of the main algorithm used by `flipper` has some important differences. These stem from the fact that the constant K of Theorem 2.10 is unknown. Therefore to decide whether a given lamination is filling `flipper` computes the splitting sequence $T_0, T_1, \ldots$ until it finds a pair which are projectively equal. Hence `flipper` will actually perform $t + m$ splits, although this is still $O(\ell(h)^2)$. As the total number of steps needed is unknown at the beginning, `flipper` records the weights of T_i algebraically as integer linear combinations of the weights of T_0. It is possible that the sequence of splits continues for so long that the precision to which these weights are known no longer uniquely determines them. In this case T_0 is recalculated to higher precision using an iterative technique and the computation can continue.

An implementation of the solution to the conjugacy problem in the pseudo-Anosov case is also included in `flipper`.

References

[1] Ian Agol. Ideal triangulations of pseudo-Anosov mapping tori. In *Topology and geometry in dimension three*, volume 560 of *Contemp. Math.*, pages 1–17. Amer. Math. Soc., Providence, RI, 2011. [5, 6, 7, 21]

[2] Ian Agol, Joel Hass, and William Thurston. The computational complexity of knot genus and spanning area. *Trans. Amer. Math. Soc.*, 358(9):3821–3850, 2006. [19, 22]

[3] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in real algebraic geometry*, volume 10 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin, second edition, 2006. [11]

[4] M. C. Bell. Deciding reducibility of mapping classes is in **NP**. *ArXiv e-prints*, March 2014. [2, 4, 20, 21, 22]

[5] Mark Bell. flipper (computer software). `https://bitbucket.org/Mark_Bell/flipper/`, 2013–2016. Version 0.9.9. [24]

[6] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999. [3]

[7] Andrew J. Casson and Steven A. Bleiler. *Automorphisms of surfaces after Nielsen and Thurston*, volume 9 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1988. [1]

[8] Henri Cohen. *A course in computational algebraic number theory*, volume 138 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, 1993. [11]

[9] Jeff Erickson and Amir Nayyeri. Tracing compressed curves in triangulated surfaces. *Discrete Comput. Geom.*, 49(4):823–863, 2013. [19, 21, 22]

[10] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012. [2, 17, 18, 22]

[11] Albert Fathi, François Laudenbach, and Valentin Poénaru. *Thurston's work on surfaces*, volume 48 of *Mathematical Notes*. Princeton University Press, Princeton, NJ, 2012. Translated from the 1979 French original by Djun M. Kim and Dan Margalit. [1]

[12] Donald E. Knuth. *The art of computer programming. Vol. 2*. Addison-Wesley, Reading, MA, 1998. Seminumerical algorithms, Third edition [of MR0286318]. [16]

[13] Thomas Koberda and Johanna Mangahas. An effective algebraic detection of the Nielsen-Thurston classification of mapping classes. *J. Topol. Anal.*, 7(1):1–21, 2015. [2, 22]

[14] Yukio Matsumoto and José María Montesinos-Amilibia. Pseudo-periodic homeomorphisms and degeneration of Riemann surfaces. *Bull. Amer. Math. Soc. (N.S.)*, 30(1):70–75, 1994. [22]

[15] Lee Mosher. Tiling the projective foliation space of a punctured surface. *Trans. Amer. Math. Soc.*, 306(1):1–70, 1988. [3]

[16] Lee Mosher. Mapping class groups are automatic. *Ann. of Math. (2)*, 142(2):303–384, 1995. [2]

[17] Lee Mosher. Train track expansions of measured foliations. Available at `http://andromeda.rutgers.edu/~mosher/`, 2003. Preprint. [24]

[18] Lee Mosher. Mapping class groups. Available at `http://andromeda.rutgers.edu/~mosher/`, 2007. Lecture notes. [18, 22, 23, 24]

[19] R. C. Penner and J. L. Harer. *Combinatorics of train tracks*, volume 125 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1992. [1, 4, 5, 8]

[20] Marcus Schaefer, Eric Sedgwick, and Daniel Štefankovič. Algorithms for normal curves and surfaces. In *Computing and combinatorics*, volume 2387 of *Lecture Notes in Comput. Sci.*, pages 370–380. Springer, Berlin, 2002. [19, 22]

[21] Marcus Schaefer, Eric Sedgwick, and Daniel Stefankovic. Computing Dehn twists and geometric intersection numbers in polynomial time, 2007. [23]

[22] Jing Tao. Linearly bounded conjugator property for mapping class groups. *Geom. Funct. Anal.*, 23(1):415–466, 2013. [2, 7]

[23] Michel Waldschmidt. *Diophantine approximation on linear algebraic groups*, volume 326 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2000. Transcendence properties of the exponential function in several variables. [11]